

On the Power of Truncated SVD for General High-rank Matrix Estimation Problems

Simon S. Du*, Yining Wang†, and Aarti Singh‡

Machine Learning Department, Carnegie Mellon University, USA

February 23, 2017

Abstract

We show that given an estimate $\hat{\mathbf{A}}$ that is close to a general high-rank positive semi-definite (PSD) matrix \mathbf{A} in spectral norm (i.e., $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$), the simple truncated SVD of $\hat{\mathbf{A}}$ produces a *multiplicative* approximation of \mathbf{A} in *Frobenius* norm. This observation leads to many interesting results on general high-rank matrix estimation problems, which we briefly summarize below (\mathbf{A} is an $n \times n$ high-rank PSD matrix and \mathbf{A}_k is the best rank- k approximation of \mathbf{A}):

- **High-rank matrix completion:** By observing $\Omega(\frac{n \max\{\epsilon^{-4}, k^2\} \mu_0^2 \|\mathbf{A}\|_F^2 \log n}{\sigma_{k+1}(\mathbf{A})^2})$ elements of \mathbf{A} where $\sigma_{k+1}(\mathbf{A})$ is the $(k+1)$ -th singular value of \mathbf{A} and μ_0 is the incoherence, the truncated SVD on a zero-filled matrix satisfies $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + O(\epsilon)) \|\mathbf{A} - \mathbf{A}_k\|_F$ with high probability. This improves results in [Har14, HW14] whose sample complexity is not gap-free, [NW12] which assumes \mathbf{A} has small Schatten- q norm, and [ZYJZ15, KLT11] which have an extra $O(\sqrt{n})$ multiplicative factor in the error bound.
- **High-rank matrix de-noising:** Let $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{E}$ where \mathbf{E} is a Gaussian random noise matrix with zero mean and ν^2/n variance on each entry. Then the truncated SVD of $\hat{\mathbf{A}}$ satisfies $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + O(\sqrt{\nu/\sigma_{k+1}(\mathbf{A})})) \|\mathbf{A} - \mathbf{A}_k\|_F + O(\sqrt{k\nu})$. This generalizes many of existing work on matrix denoising [DG⁺14, DGM13, GD14] where \mathbf{A} is assumed to be exact or nearly low-rank.
- **Low-rank Estimation of high-dimensional covariance:** Given N i.i.d. samples $X_1, \dots, X_N \sim \mathcal{N}_n(\mathbf{0}, \mathbf{A})$, can we estimate \mathbf{A} with a relative-error Frobenius norm bound? We show that if $N = \Omega(n \max\{\epsilon^{-4}, k^2\} \gamma_k(\mathbf{A})^2 \log N)$ for $\gamma_k(\mathbf{A}) = \sigma_1(\mathbf{A})/\sigma_{k+1}(\mathbf{A})$, then $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + O(\epsilon)) \|\mathbf{A} - \mathbf{A}_k\|_F$ with high probability, where $\hat{\mathbf{A}} = \frac{1}{N} \sum_{i=1}^N X_i X_i^\top$ is the sample covariance. This generalizes previous results [BX⁺15] which applies to covariance \mathbf{A} with small effective rank only, at the cost of worse dependency on ϵ .

1 Introduction

Let \mathbf{A} be an unknown general high-rank $n \times n$ PSD data matrix that one wishes to estimate. In many machine learning applications, though \mathbf{A} is unknown, it are relatively easy to obtain a crude estimate $\hat{\mathbf{A}}$ that is close to \mathbf{A} in spectral norm (i.e., $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$). For example, in matrix completion a simple procedure that fills all unobserved entries with 0 and re-scales observed entries produces an estimate that is consistent in

*ssdu@cs.cmu.edu

†yiningwa@cs.cmu.edu

‡aarti@cs.cmu.edu

spectral norm (assuming the matrix satisfies a spikeness condition, standard assumption in matrix completion literature). In matrix de-noising, an observation that is corrupted by Gaussian noise is close to the underlying signal, because standard Gaussian noise is isotropic and has small spectral norm. In covariance estimation, the sample covariance in low-dimensional settings is close to the population covariance in spectral norm under mild conditions [BX⁺15].

However, in most such applications it is not sufficient to settle for a spectral norm approximation. For example, in recommendation systems (an application of matrix completion) the zero-filled re-scaled rating matrix is close to the ground truth in spectral norm, but it is an absurd estimator because most of the estimated ratings are zero. It is hence mandatory to require a more stringent measure of performance. One commonly used measure is the *Frobenius norm* of the estimation error $\|\hat{\mathbf{A}} - \mathbf{A}\|_F$, which ensures that (on average) the estimate is close to the ground truth in an element-wise sense. A spectral norm approximation $\hat{\mathbf{A}}$ is in general *not* a good estimate under Frobenius norm, because in high-rank scenarios $\|\hat{\mathbf{A}} - \mathbf{A}\|_F$ can be \sqrt{n} times larger than $\|\hat{\mathbf{A}} - \mathbf{A}\|_2$.

In this paper, we show that in many cases a powerful multiplicative low-rank approximation in Frobenius norm can be obtained by applying a simple truncated SVD procedure on a crude, easy-to-find spectral norm approximate. In particular, given the spectral norm approximation condition $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$, the top- k SVD of $\hat{\mathbf{A}}_k$ of $\hat{\mathbf{A}}$ multiplicatively approximates \mathbf{A} in Frobenius norm; that is, $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq C(k, \delta, \sigma_{k+1}(\mathbf{A}))\|\mathbf{A} - \mathbf{A}_k\|_F$, where \mathbf{A}_k is the best rank- k approximation of \mathbf{A} in Frobenius and spectral norm. To our knowledge, the best existing result under the assumption $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$ is due to [AM07], who showed that $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq \|\mathbf{A} - \mathbf{A}_k\|_F + \sqrt{k}\delta + 2k^{1/4}\sqrt{\delta}\|\mathbf{A}_k\|_F$, which depends on $\|\mathbf{A}_k\|_F$ and is not multiplicative in $\|\mathbf{A} - \mathbf{A}_k\|_F$.

Below we summarize applications in several matrix estimation problems, and comment on the proof technique that we employ to prove the main result.

High-rank matrix completion Matrix completion is the problem of (approximately) recovering a data matrix from very few observed entries. It has wide applications in machine learning, especially in online recommendation systems. Most existing work on matrix completion assumes the data matrix is *exactly* low-rank [CR12, SL16, JNS13]. [CP10, KMO10] studied the problem of recovering a low-rank matrix corrupted by stochastic noise; [CXCS16] considered sparse column corruption. All of the aforementioned work assumes that the ground-truth data matrix is exactly low-rank, which is rarely true in practice.

[NW12] derived minimax rates of estimation error when the spectrum of the data matrix lies in an ℓ_q ball. [ZYJZ15, KLT11] derived oracle inequalities for general matrix completion; however their error bound has an additional $O(\sqrt{n})$ multiplicative factor. These results also require solving computationally expensive nuclear-norm penalized optimization problems whereas our method only requires solving a single truncated singular value decomposition. [HW14] used a “soft-deflation” technique to remove condition number dependency in the sample complexity; however, their error bound for general high-rank matrix completion is additive and depends on the “consecutive” spectral gap $(\sigma_k(\mathbf{A}) - \sigma_{k+1}(\mathbf{A}))$, which can be small in practical settings [BDWY16]. [EBN12] considered high-rank matrix completion with additional union-of-subspace structures.

In this paper, we show that if the $n \times n$ data matrix \mathbf{A} satisfies μ_0 -spikeness condition,¹ then for any $\epsilon \in (0, 1)$, the truncated SVD of zero-filled matrix $\hat{\mathbf{A}}_k$ satisfies $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + O(\epsilon))\|\mathbf{A} - \mathbf{A}_k\|_F$ if the sample complexity is lower bounded by $\Omega(\frac{n \max\{\epsilon^{-4}, k^2\} \mu_0^2 \|\mathbf{A}\|_F^2 \log n}{\sigma_{k+1}(\mathbf{A})^2})$, which can be further simplified to $\Omega(\mu_0^2 \max\{\epsilon^{-4}, k^2\} \gamma_k(\mathbf{A})^2 \cdot nr_s(\mathbf{A}) \log n)$, where $\gamma_k(\mathbf{A}) = \sigma_1(\mathbf{A})/\sigma_{k+1}(\mathbf{A})$ is the k th-order condition number and $r_s(\mathbf{A}) = \|\mathbf{A}\|_F^2 / \|\mathbf{A}\|_2^2 \leq \text{rank}(\mathbf{A})$ is the *stable rank* of \mathbf{A} . Compared to existing work, our error bound is multiplicative, gap-free, and the estimator is computationally efficient.²

High-rank matrix de-noising Let $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{E}$ be a noisy observation of \mathbf{A} , where \mathbf{E} is a PSD Gaussian noise matrix with zero mean and ν^2/n variance on each entry. By simple concentration results we have

¹ $n\|\mathbf{A}\|_{\max} \leq \mu_0\|\mathbf{A}\|_F$; see also Definition 2.1.

² We remark that our analysis does *not*, however, apply to exact rank- k matrix where $\sigma_{k+1} = 0$. This is because for exact rank- k matrix a bound of the form $(1 + O(\epsilon))\|\mathbf{A} - \mathbf{A}_k\|_F$ requires *exact recovery* of \mathbf{A} , which truncated SVD cannot achieve, but it is achievable when combining with projected gradient descent [JN14].

$\|\hat{\mathbf{A}} - \mathbf{A}\|_2 = \nu$ with high probability; however, $\hat{\mathbf{A}}$ is in general not a good approximation of \mathbf{A} in Frobenius norm when \mathbf{A} is high-rank. More specifically, $\|\hat{\mathbf{A}} - \mathbf{A}\|_F$ can be as large as $\sqrt{n}\nu$.

Applying our main result, we show that if $\nu < \sigma_{k+1}(\mathbf{A})$ for some $k \ll n$, then the top- k SVD $\hat{\mathbf{A}}_k$ of $\hat{\mathbf{A}}$ satisfies $\|\hat{\mathbf{A}}_k - \mathbf{A}_k\|_F \leq (1 + O(\sqrt{\nu/\sigma_{k+1}(\mathbf{A})}))\|\mathbf{A} - \mathbf{A}_k\|_F + \sqrt{k}\nu$. This suggests a form of bias-variance decomposition as larger rank threshold k induces smaller bias $\|\mathbf{A} - \mathbf{A}_k\|_F$ but larger variance $k\nu^2$. Our results generalize existing work on matrix de-noising [DG⁺14, DGM13, GD14], which focus primarily on exact low-rank \mathbf{A} .

Low-rank estimation of high-dimensional covariance The (Gaussian) covariance estimation problem asks to estimate an $n \times n$ PSD covariance matrix \mathbf{A} , either in spectral or Frobenius norm, from N i.i.d. samples $X_1, \dots, X_N \sim \mathcal{N}(\mathbf{0}, \mathbf{A})$. The *high-dimensional* regime of covariance estimation, in which $N \approx n$ or even $N \gg n$, has attracted enormous interest in the mathematical statistics literature [CZZ⁺10, CZ⁺12, CRZ13, CRZ⁺16]. While most existing work focus on sparse or banded covariance matrices, the setting where \mathbf{A} has certain low-rank structure has seen rising interest recently [BX⁺15, KS11]. In particular, [BX⁺15] shows that if $n = O(N^\beta)$ for some $\beta \geq 0$ then the sample covariance estimator $\hat{\mathbf{A}} = \frac{1}{N} \sum_{i=1}^N X_i X_i^\top$ satisfies

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_F = O_{\mathbb{P}} \left(\|\mathbf{A}\|_2 r_e(\mathbf{A}) \sqrt{\frac{\log N}{N}} \right), \quad (1)$$

where $r_e(\mathbf{A}) = \text{tr}(\mathbf{A})/\|\mathbf{A}\|_2 \leq \text{rank}(\mathbf{A})$ is the *effective rank* of \mathbf{A} . For high-rank matrices where $r_e(\mathbf{A}) \approx n$, Eq. (1) requires $N = \Omega(n^2 \log n)$ to approximate \mathbf{A} consistently in Frobenius norm.

In this paper we consider a reduced-rank estimator $\hat{\mathbf{A}}_k$ and show that, if $\frac{r_e(\mathbf{A}) \max\{\epsilon^{-4}, k^2\} \gamma_k(\mathbf{A})^2 \log N}{N} \leq c$ for some small universal constant $c > 0$, then $\|\hat{\mathbf{A}}_k - \mathbf{A}_k\|_F$ admits a relative Frobenius-norm error bound $(1 + O(\epsilon))\|\mathbf{A} - \mathbf{A}_k\|_F$ with high probability. Our result allows reasonable approximation of \mathbf{A} in Frobenius norm under the regime of $N = \Omega(n \text{poly}(k) \log n)$ if $\gamma_k = O(\text{poly}(k))$, which is significantly more flexible than $N = \Omega(n^2 \log n)$, though the dependency of ϵ is worse than [BX⁺15]. The error bound is also agnostic in nature, making no assumption on the actual or effective rank of \mathbf{A} .

Proof techniques Because both $\hat{\mathbf{A}}_k$ and \mathbf{A}_k are low-rank, $\|\hat{\mathbf{A}}_k - \mathbf{A}_k\|_F$ is upper bounded by an $O(\sqrt{k})$ factor of $\|\hat{\mathbf{A}}_k - \mathbf{A}_k\|_2$. From the condition that $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$, a straightforward approach to upper bound $\|\hat{\mathbf{A}}_k - \mathbf{A}_k\|_2$ is to consider the decomposition $\|\hat{\mathbf{A}}_k - \mathbf{A}_k\|_2 \leq \|\hat{\mathbf{A}} - \mathbf{A}\|_2 + 2\|\mathbf{U}_k \mathbf{U}_k^\top - \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top\|_2 \|\hat{\mathbf{A}}_k\|_2$, where $\mathbf{U}_k \mathbf{U}_k^\top$ and $\hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top$ are projection operators onto the top- k eigenspaces of \mathbf{A} and $\hat{\mathbf{A}}$, respectively. Such a naive approach, however, has two major disadvantages. First, the upper bound depends on $\|\hat{\mathbf{A}}_k\|_2$, which is additive and may be much larger than $\|\hat{\mathbf{A}} - \mathbf{A}\|_2$. Perhaps more importantly, the quantity $\|\mathbf{U}_k \mathbf{U}_k^\top - \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top\|_2$ depends on the “consecutive” spectral gap $(\sigma_k(\mathbf{A}) - \sigma_{k+1}(\mathbf{A}))$, which could be very small for large matrices.

To obtain a multiplicative error bound and remove dependency over consecutive spectral gap $(\sigma_k(\mathbf{A}) - \sigma_{k+1}(\mathbf{A}))$, the key idea is to consider an “envelope” of spectrum $\sigma_{m_1}(\mathbf{A}) \geq \sigma_k(\mathbf{A}) \geq \sigma_{m_2}(\mathbf{A})$ such that $\sigma_{m_1}(\mathbf{A}) \approx (1 + 2\epsilon)\sigma_k(\mathbf{A})$ and $\sigma_{m_2}(\mathbf{A}) \approx \sigma_k(\mathbf{A}) - 2\epsilon\sigma_{k+1}(\mathbf{A})$. This allows the usage of an asymmetric Davis-Kahan Theorem to establish gap-free matrix perturbation bounds. Similar ideas were discussed in [AZL16] to analyze shift-and-invert and inexact power method for fast SVD computation.

Notations For an $n \times n$ PSD matrix \mathbf{A} , denote $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^\top$ as its eigenvalue decomposition, where \mathbf{U} is an orthogonal matrix and $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_n)$ is a diagonal matrix, with eigenvalues sorted in descending order $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$. The spectral norm and Frobenius norm of \mathbf{A} are defined as $\|\mathbf{A}\|_2 = \sigma_1$ and $\|\mathbf{A}\|_F = \sqrt{\sigma_1^2 + \dots + \sigma_n^2}$, respectively. Suppose $\mathbf{u}_1, \dots, \mathbf{u}_n$ are eigenvectors associated with $\sigma_1, \dots, \sigma_n$. Define $\mathbf{A}_k = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{u}_i^\top = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{U}_k^\top$, $\mathbf{A}_{n-k} = \sum_{i=k+1}^n \sigma_i \mathbf{u}_i \mathbf{u}_i^\top = \mathbf{U}_{n-k} \mathbf{\Sigma}_{n-k} \mathbf{U}_{n-k}^\top$ and $\mathbf{A}_{m_1:m_2} = \sum_{i=m_1+1}^{m_2} \sigma_i \mathbf{u}_i \mathbf{u}_i^\top = \mathbf{U}_{m_1:m_2} \mathbf{\Sigma}_{m_1:m_2} \mathbf{U}_{m_1:m_2}^\top$. For a tall matrix $\mathbf{U} \in \mathbb{R}^{n \times k}$, we use $\mathcal{U} = \text{Range}(\mathbf{U})$ to denote the linear subspace spanned by the columns of \mathbf{U} . For two linear subspaces \mathcal{U} and \mathcal{V} , we write $\mathcal{W} = \mathcal{U} \oplus \mathcal{V}$ if $\mathcal{U} \cap \mathcal{V} = \{\mathbf{0}\}$ and $\mathcal{W} = \{\mathbf{u} + \mathbf{v} : \mathbf{u} \in \mathcal{U}, \mathbf{v} \in \mathcal{V}\}$. For a sequence of random variables

$\{X_n\}_{n=1}^\infty$ and real-valued function $f : \mathbb{N} \rightarrow \mathbb{R}$, we say $X_n = O_{\mathbb{P}}(f(n))$ if for any $\epsilon > 0$, there exists $N \in \mathbb{N}$ and $C > 0$ such that $\Pr[|X_n| \geq C \cdot |f(n)|] \leq \epsilon$ for all $n \geq N$.

2 Multiplicative Frobenius-norm approximation and applications

We first state our main result, which shows that truncated SVD on a weak estimator with small approximation error in spectral norm leads to a strong estimator with *multiplicative* Frobenius-norm error bound. We remark that truncated SVD in general has time complexity

$$O\left(\min\left\{n^2k, \text{nnz}(\tilde{\mathbf{A}}) + n\text{poly}(k)\right\}\right),$$

where $\text{nnz}(\tilde{\mathbf{A}})$ is the number of non-zero entries in $\tilde{\mathbf{A}}$, and the time complexity is at most linear in matrix sizes when k is small. We refer readers to [AZL16] for details. The proof of the main theorem is placed in Sec. 3.

Theorem 2.1. *Suppose \mathbf{A} is an $n \times n$ PSD matrix with eigenvalues $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_n(\mathbf{A}) \geq 0$, and a symmetric matrix $\hat{\mathbf{A}}$ satisfies $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta = \epsilon^2 \sigma_{k+1}(\mathbf{A})$ for some $\epsilon \in (0, 1/4]$. Let \mathbf{A}_k and $\hat{\mathbf{A}}_k$ be the best rank- k approximations of \mathbf{A} and $\hat{\mathbf{A}}$. Then*

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + 32\epsilon)\|\mathbf{A} - \mathbf{A}_k\|_F + 102\sqrt{2k}\epsilon^2\|\mathbf{A} - \mathbf{A}_k\|_2. \quad (2)$$

Remark: Note when $\epsilon = O(1/\sqrt{k})$ we obtain an $(1 + O(\epsilon))$ error bound.

To our knowledge, the best existing bound for $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F$ assuming $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$ is due to [AM07], who showed that

$$\begin{aligned} \|\hat{\mathbf{A}}_k - \mathbf{A}\|_F &\leq \|\mathbf{A} - \mathbf{A}_k\|_F + \|(\hat{\mathbf{A}} - \mathbf{A})_k\|_F + 2\sqrt{\|(\hat{\mathbf{A}} - \mathbf{A})_k\|_F \|\mathbf{A}_k\|_F} \\ &\leq \|\mathbf{A} - \mathbf{A}_k\|_F + \sqrt{k}\delta\|\mathbf{A} - \mathbf{A}_k\|_2 + 2k^{1/4}\sqrt{\delta}\sqrt{\|\mathbf{A}_k\|_F}. \end{aligned} \quad (3)$$

Compared to Theorem 2.1, Eq. (3) is not relative because the third term $2k^{1/4}\sqrt{\|\mathbf{A}_k\|_F}$ depends on the k largest eigenvalues of \mathbf{A} , which could be much larger than the remainder term $\|\mathbf{A} - \mathbf{A}_k\|_F$. In contrast, Theorem 2.1, together with Remark 2, shows that $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F$ could be upper bounded by a small factor multiplied with the remainder term $\|\mathbf{A} - \mathbf{A}_k\|_F$.

Before we proceed to the applications and proof of Theorem 2.1, we first list several examples of \mathbf{A} with classical distribution of eigenvalues and discuss how Theorem 2.1 could be applied to obtain good Frobenius-norm approximations of \mathbf{A} . We begin with the case where eigenvalues of \mathbf{A} have a polynomial decay rate (i.e., power law). Such matrices are ubiquitous in practice [LWS15].

Corollary 2.1 (Power-law spectral decay). *Suppose $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$ for some $\delta \in (0, 1/2]$ and $\sigma_j(\mathbf{A}) = j^{-\beta}$ for some $\beta > 1/2$. Set $k = \lfloor \min\{C_1\delta^{-1/\beta}, n\} - 1 \rfloor$. If $k \geq 1$ then*

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq C'_1 \cdot \max\left\{\delta^{\frac{2\beta-1}{2\beta}}, n^{-\frac{2\beta-1}{2\beta}}\right\},$$

where $C_1, C'_1 > 0$ are constants that only depend on β .

We remark that the assumption $\sigma_j(\mathbf{A}) = j^{-\beta}$ implies that the eigenvalues lie in an ℓ_q ball for $q = 1/\beta$; that is, $\sum_{j=1}^n \sigma_j(\mathbf{A})^q = O(1)$. The error bound in Corollary 2.1 matches the minimax rate (derived in [NW12]) for matrix completion when the spectrum is constrained in an ℓ_q ball, by replacing δ with $\sqrt{n/N}$ where N is the number of observed entries.

Next, we consider the case where eigenvalues satisfy a faster decay rate.

Corollary 2.2 (Exponential spectral decay). *Suppose $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \delta$ for some $\delta \in (0, e^{-16})$ and $\sigma_j(\mathbf{A}) = \exp\{-cj\}$ for some $c > 0$. Set $k = \lfloor \min\{c^{-1} \log(1/\delta) - c^{-1} \log \log(1/\delta), n\} - 1 \rfloor$. If $k \geq 1$ then*

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq C'_2 \cdot \max \left\{ \delta \sqrt{\log(1/\delta)^3}, n^{1/2} \exp(-cn) \right\},$$

where $C'_2 > 0$ is a constant that only depends on c .

Both corollaries are proved in the appendix. The error bounds in both Corollaries 2.1 and 2.2 are significantly better than the trivial estimate $\hat{\mathbf{A}}$, which satisfies $\|\hat{\mathbf{A}} - \mathbf{A}\|_F \leq n^{1/2}\delta$. We also remark that the bound in Corollary 2.1 cannot be obtained by a direct application of the weaker bound Eq. (3), which yields a $\delta^{\frac{\beta}{2\beta-1}}$ bound.

We next state results that are consequences of Theorem 2.1 in several matrix estimation problems.

2.1 High-rank matrix completion

Suppose \mathbf{A} is a high-rank $n \times n$ PSD matrix that satisfies μ_0 -spikeness condition defined as follows:

Definition 2.1 (Spikeness condition). *An $n \times n$ PSD matrix \mathbf{A} satisfies μ_0 -spikeness condition if $n\|\mathbf{A}\|_{\max} \leq \mu_0\|\mathbf{A}\|_F$, where $\|\mathbf{A}\|_{\max} = \max_{1 \leq i, j \leq n} |\mathbf{A}_{ij}|$ is the max-norm of \mathbf{A} .*

Spikeness condition makes uniform sampling of matrix entries powerful in matrix completion problems. If \mathbf{A} is exactly low rank, the spikeness condition is implied by an upper bound on $\max_{1 \leq i \leq n} \|\mathbf{e}_i^\top \mathbf{U}_k\|_2$, which is the standard incoherence assumption on the top- k space of \mathbf{A} [CR12]. For general high-rank \mathbf{A} , the spikeness condition is implied by a more restrictive incoherence condition that imposes an upper bound on $\max_{1 \leq i \leq n} \|\mathbf{e}_i^\top \mathbf{U}_{n-k}\|_2$ and $\|\mathbf{A}_{n-k}\|_{\max}$, which are assumptions adopted in [HW14].

Suppose $\hat{\mathbf{A}}$ is a symmetric re-scaled zero-filled matrix of observed entries. That is,

$$[\hat{\mathbf{A}}]_{ij} = \begin{cases} \mathbf{A}_{ij}/p, & \text{with probability } p; \\ 0, & \text{with probability } 1 - p; \end{cases} \quad \forall 1 \leq i \leq j \leq n. \quad (4)$$

Here $p \in (0, 1)$ is a parameter that controls the probability of observing a particular entry in \mathbf{A} , corresponding to a sample complexity of $O(n^2/p)$. Note that both $\hat{\mathbf{A}}$ and \mathbf{A} are symmetric so we only specify the upper triangle of $\hat{\mathbf{A}}$. By a simple application of matrix Bernstein inequality [MJC⁺14], one can show $\hat{\mathbf{A}}$ is close to \mathbf{A} in spectral norm when \mathbf{A} satisfies μ_0 -spikeness. Here we cite a lemma from [Har14] to formally establish this observation:

Lemma 2.1 ([Har14], Lemma A.3). *Under the model of Eq. (4), for $u > 0$ it holds that*

$$\Pr \left[\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \geq u \right] \leq 2n \exp \left\{ - \frac{u^2/2}{(1/p - 1) \max_{1 \leq i \leq n} \|\mathbf{e}_i^\top \mathbf{A}\|_2^2 + \frac{u}{3}(1/p - 1)\|\mathbf{A}\|_\infty} \right\}.$$

Corollary 2.3. *Under the model of Eq. (4) and μ_0 -spikeness condition of \mathbf{A} , for $t \in (0, 1)$ it holds with probability at least $1 - t$ that*

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq O \left(\max \left\{ \sqrt{\frac{\mu_0^2 \|\mathbf{A}\|_F^2 \log(n/t)}{np}}, \frac{\mu_0 \|\mathbf{A}\|_F \log(n/t)}{np} \right\} \right).$$

Let $\hat{\mathbf{A}}_k$ be the best rank- k approximation of $\hat{\mathbf{A}}$ in Frobenius/spectral norm. Applying Theorem 2.1 we obtain the following result:

Theorem 2.2. Fix $t \in (0, 1)$. Then with probability $1 - t$ we have

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq O(\sqrt{k}) \cdot \|\mathbf{A} - \mathbf{A}_k\|_F \quad \text{if} \quad p = \Omega\left(\frac{\mu_0^2 \|\mathbf{A}\|_F^2 \log(n/t)}{n\sigma_{k+1}(\mathbf{A})^2}\right). \quad (5)$$

Furthermore, for fixed $\epsilon \in (0, 1/4]$, with probability $1 - t$ we have

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + O(\epsilon)) \|\mathbf{A} - \mathbf{A}_k\|_F \quad \text{if} \quad p = \Omega\left(\frac{\mu_0^2 \max\{\epsilon^{-4}, k^2\} \|\mathbf{A}_0\|_F^2 \log(n/t)}{n\sigma_{k+1}(\mathbf{A})^2}\right). \quad (6)$$

As a remark, because $\mu_0 \geq 1$ and $\|\mathbf{A}\|_F/\sigma_{k+1}(\mathbf{A}) \geq \sqrt{k}$ always hold, the sample complexity in both Eq. (5,6) is lower bounded by $\Omega(nk \log n)$, the typical sample complexity in noiseless matrix completion. In the case of high rank \mathbf{A} , the results in Theorem 2.2 are the strongest when \mathbf{A} has small *stable rank* $r_s(\mathbf{A}) = \|\mathbf{A}\|_F^2/\|\mathbf{A}\|_2^2$ and the top- k condition number $\gamma_k(\mathbf{A}) = \sigma_1(\mathbf{A})/\sigma_{k+1}(\mathbf{A})$ is not too large. For example, if \mathbf{A} has stable rank $r_s(\mathbf{A}) = r$ then $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F$ has an $O(\sqrt{k})$ multiplicative error bound with sample complexity $\Omega(\mu_0^2 \gamma_k(\mathbf{A})^2 \cdot nr \log n)$; or an $(1 + O(\epsilon))$ relative error bound with sample complexity $\Omega(\mu_0^2 \max\{\epsilon^{-4}, k^2\} \gamma_k(\mathbf{A})^2 \cdot nr \log n)$.

2.2 High-rank matrix de-noising

Let \mathbf{A} be an $n \times n$ PSD signal matrix and \mathbf{E} a symmetric random Gaussian matrix with zero mean and ν^2/n variance. That is, $\mathbf{E}_{ij} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \nu^2/n)$ for $1 \leq i \leq j \leq n$ and $\mathbf{E}_{ij} = \mathbf{E}_{ji}$. Define $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{E}$. The matrix de-noising problem is then to recover the signal matrix \mathbf{A} from noisy observations $\hat{\mathbf{A}}$. We refer the readers to [GD14] for a list of references that shows the ubiquitous application of matrix de-noising in scientific fields.

It is well-known by concentration results of Gaussian random matrices, that $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 = \|\mathbf{E}\|_2 = O_{\mathbb{P}}(\nu)$. Let $\hat{\mathbf{A}}_k$ be the best rank- k approximation of $\hat{\mathbf{A}}$ in Frobenius/spectral norm. Applying Theorem 2.1, we immediately have the following result:

Theorem 2.3. There exists an absolute constant $c > 0$ such that, if $\nu < c \cdot \sigma_{k+1}(\mathbf{A})$ for some $1 \leq k < n$, then with probability at least 0.8 we have that

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq \left(1 + O\left(\sqrt{\frac{\nu}{\sigma_{k+1}(\mathbf{A})}}\right)\right) \|\mathbf{A} - \mathbf{A}_k\|_F + O(\sqrt{k}\nu). \quad (7)$$

Eq. (7) can be understood from a classical bias-variance tradeoff perspective: the first $(1 + O(\sqrt{\nu/\sigma_{k+1}(\mathbf{A})}))\|\mathbf{A} - \mathbf{A}_k\|_F$ acts as a bias term, which decreases as we increase cut-off rank k , corresponding to a more complicated model; on the other hand, the second $O(\sqrt{k}\nu)$ term acts as the (square root of) variance, which does not depend on the signal \mathbf{A} and increases with k .

2.3 Low-rank estimation of high-dimensional covariance

Suppose \mathbf{A} is an $n \times n$ PSD matrix and X_1, \dots, X_N are i.i.d. samples drawn from the multivariate Gaussian distribution $\mathcal{N}_n(\mathbf{0}, \mathbf{A})$. The question is to estimate \mathbf{A} from samples X_1, \dots, X_N . A common estimator is the *sample covariance* $\hat{\mathbf{A}} = \frac{1}{N} \sum_{i=1}^N X_i X_i^\top$. While in low-dimensional regimes (i.e., n fixed and $N \rightarrow \infty$) the asymptotic efficiency of $\hat{\mathbf{A}}$ is obvious (cf. [VdV00]), its statistical power in high-dimensional regimes where n and N are comparable are highly non-trivial. Below we cite results in [BX⁺15] for estimation error $\|\hat{\mathbf{A}} - \mathbf{A}\|_\xi$, $\xi = 2/F$ when n is not too large compared to N :

Lemma 2.2 ([BX⁺15]). Suppose $n = O(N^\beta)$ for some $\beta \geq 0$ and let $r_e(\mathbf{A}) = \text{tr}(\mathbf{A})/\|\mathbf{A}\|_2$ denote the effective rank of the covariance \mathbf{A} . Then the sample covariance $\hat{\mathbf{A}} = \frac{1}{N} \sum_{i=1}^N X_i X_i^\top$ satisfies

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_F = O_{\mathbb{P}} \left(\|\mathbf{A}\|_2 r_e(\mathbf{A}) \sqrt{\frac{\log N}{N}} \right) \quad (8)$$

and

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_2 = O_{\mathbb{P}} \left(\|\mathbf{A}\|_2 \max \left\{ \sqrt{\frac{r_e(\mathbf{A}) \log(Nn)}{N}}, \frac{r_e(\mathbf{A}) \log(Nn)}{N} \right\} \right). \quad (9)$$

Let $\hat{\mathbf{A}}_k$ be the best rank- k approximation of $\hat{\mathbf{A}}$ in Frobenius/spectral norm. Applying Theorem 2.1 together with Eq. (9), we immediately arrive at the following theorem that gives relative-error Frobenius norm bound for $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F$:

Theorem 2.4. Fix $\epsilon \in (0, 1/4]$ and $1 \leq k < n$. Recall that $r_e(\mathbf{A}) = \text{tr}(\mathbf{A})/\|\mathbf{A}\|_2$ and $\gamma_k(\mathbf{A}) = \sigma_1(\mathbf{A})/\sigma_{k+1}(\mathbf{A})$. There exists a universal constant $c > 0$ such that, if

$$\frac{r_e(\mathbf{A}) \max\{\epsilon^{-4}, k^2\} \gamma_k(\mathbf{A})^2 \log(N)}{N} \leq c$$

then with probability at least 0.8,

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq (1 + O(\epsilon)) \|\mathbf{A} - \mathbf{A}_k\|_F.$$

Theorem 2.4 shows that it is possible to obtain a reasonable Frobenius-norm approximation of $\hat{\mathbf{A}}$ by truncated SVD in the asymptotic regime of $N = \Omega(r_e(\mathbf{A}) \text{poly}(k) \log N)$, which is much more flexible than Eq. (8) that requires $N = \Omega(r_e(\mathbf{A})^2 \log N)$.

3 Proof of Theorem 2.1

The key idea in the proof of Theorem 2.1 is to find an “envelope” $m_1 \leq k \leq m_2$ in the spectrum of \mathbf{A} surrounding k , such that the eigenvalues within the envelope are relatively close. Define

$$\begin{aligned} m_1 &= \text{argmax}_{0 \leq j \leq k} \{\sigma_j(\mathbf{A}) \geq (1 + 2\epsilon)\sigma_{k+1}(\mathbf{A})\}; \\ m_2 &= \text{argmax}_{k \leq j \leq n} \{\sigma_j(\mathbf{A}) \geq \sigma_k(\mathbf{A}) - 2\epsilon\sigma_{k+1}(\mathbf{A})\}, \end{aligned}$$

where we let $\sigma_0(\mathbf{A}) = \infty$ for convenience. Let $\mathcal{U}_m, \hat{\mathcal{U}}_m$ be basis of the top m -dimensional linear subspaces of \mathbf{A} and $\hat{\mathbf{A}}$, respectively. Also denote \mathcal{U}_{n-m} and $\hat{\mathcal{U}}_{n-m}$ as basis of the orthogonal complement of \mathcal{U}_m and $\hat{\mathcal{U}}_m$.

Lemma 3.1. If $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$ for $\epsilon \in (0, 1)$ then $\|\hat{\mathbf{U}}_{n-k}^\top \mathbf{U}_{m_1}\|_2, \|\hat{\mathbf{U}}_k^\top \mathbf{U}_{n-m_2}\|_2 \leq \epsilon$.

Proof. We apply an asymmetric version of Davis-Kahan inequality (Lemma B.1), with $\mathbf{X} = \mathbf{A}$, $\mathbf{Y} = \hat{\mathbf{A}}$, $i = m_1$ and $j = k$. By Weyl’s inequality, we know that $\sigma_{k+1}(\hat{\mathbf{A}}) \leq \sigma_{k+1}(\mathbf{A}) + \|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq (1 + \epsilon^2)\sigma_{k+1}(\mathbf{A}) \leq (1 + \epsilon)\sigma_{k+1}(\mathbf{A})$. Subsequently, $\|\hat{\mathbf{U}}_{n-k}^\top \mathbf{U}_{m_1}\|_2 \leq \frac{\epsilon^2 \sigma_{k+1}(\mathbf{A})}{\sigma_{m_1}(\mathbf{A}) - (1 + \epsilon)\sigma_{k+1}(\mathbf{A})} \leq \epsilon$. Similarly, applying Lemma B.1 with $\mathbf{X} = \hat{\mathbf{A}}$, $\mathbf{Y} = \mathbf{A}$, $i = k$ and $j = m_2$ we have that $\|\hat{\mathbf{U}}_k^\top \mathbf{U}_{n-m_2}\|_2 \leq \epsilon$. \square

Let $\mathcal{U}_{m_1:m_2}$ be the linear subspace of \mathbf{A} associated with eigenvalues $\sigma_{m_1+1}(\mathbf{A}), \dots, \sigma_{m_2}(\mathbf{A})$. Intuitively, we choose a $(k - m_1)$ -dimensional linear subspace in $\mathcal{U}_{m_1:m_2}$ that is “most aligned” with the top- k subspace $\hat{\mathcal{U}}_k$ of $\hat{\mathbf{A}}$. Formally, define

$$\mathcal{W} = \text{argmax}_{\dim(\mathcal{W})=k-m_1, \mathcal{W} \in \mathcal{U}_{m_1:m_2}} \sigma_{k-m_1}(\mathbf{W}^\top \hat{\mathbf{U}}_k).$$

\mathbf{W} is then a $d \times (k - m_1)$ matrix with orthonormal columns that corresponds to a basis of \mathcal{W} . \mathcal{W} is carefully constructed so that it is closely aligned with $\hat{\mathcal{U}}_k$, yet still lies in \mathcal{U}_k . In particular, Lemma 3.2 shows that $\sin \angle(\mathcal{W}, \hat{\mathcal{U}}_k) = \|\hat{\mathbf{U}}_{n-k}^\top \mathbf{W}\|_2$ is upper bounded by ϵ .

Lemma 3.2. *If $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$ for $\epsilon \in (0, 1)$ then $\|\hat{\mathbf{U}}_{n-k}^\top \mathbf{W}\|_2 \leq \epsilon$.*

Proof. First note that $\|\hat{\mathbf{U}}_{n-k}^\top \mathbf{W}\|_2 \leq \sqrt{1 - \sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{W})^2}$ because

$$\begin{aligned} \|\hat{\mathbf{U}}_{n-k}^\top \mathbf{W}\|_2^2 &= \sup_{\|\mathbf{x}\|_2=1} \|\hat{\mathbf{U}}_{n-k}^\top \mathbf{W} \mathbf{x}\|_2^2 = \sup_{\|\mathbf{x}\|_2=1} \left\{ \|\mathbf{W} \mathbf{x}\|_2^2 - \|\hat{\mathbf{U}}_k^\top \mathbf{W} \mathbf{x}\|_2^2 \right\} \\ &\leq \sup_{\|\mathbf{x}\|_2=1} \|\mathbf{W} \mathbf{x}\|_2^2 - \inf_{\|\mathbf{x}\|_2=1} \|\hat{\mathbf{U}}_k^\top \mathbf{W} \mathbf{x}\|_2^2 = 1 - \sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{W})^2. \end{aligned}$$

Subsequently, it suffices to prove that $\sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{W}) \geq \sqrt{1 - \epsilon^2}$. By Weyl's monotonicity theorem (Lemma B.4), we have that

$$\sigma_k(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_2}) \leq \sigma_{m_1+1}(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_1}) + \sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_1:m_2}).$$

In addition, $\sigma_{m_1+1}(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_1}) = 0$ because $\text{rank}(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_1}) \leq m_1$ and $\sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_1:m_2}) = \sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{W})$ because of the definition of \mathbf{W} . Subsequently,

$$\begin{aligned} \sigma_{k-m_1}(\hat{\mathbf{U}}_k^\top \mathbf{W})^2 &\geq \sigma_k(\hat{\mathbf{U}}_k^\top \mathbf{U}_{m_2})^2 = \inf_{\|\mathbf{x}\|_2=1} \|\mathbf{U}_{m_2}^\top \hat{\mathbf{U}}_k \mathbf{x}\|_2^2 = \inf_{\|\mathbf{x}\|_2=1} \left\{ \|\hat{\mathbf{U}}_k \mathbf{x}\|_2^2 - \|\mathbf{U}_{n-m_2}^\top \hat{\mathbf{U}}_k \mathbf{x}\|_2^2 \right\} \\ &\geq \inf_{\|\mathbf{x}\|_2=1} \left\{ \|\hat{\mathbf{U}}_k \mathbf{x}\|_2^2 \right\} - \sup_{\|\mathbf{x}\|_2=1} \left\{ \|\mathbf{U}_{n-m_2}^\top \hat{\mathbf{U}}_k \mathbf{x}\|_2^2 \right\} \geq 1 - \epsilon^2. \end{aligned}$$

Here in the last inequality we invoke Lemma 3.1. The proof is then complete. \square

Define

$$\tilde{\mathbf{A}} = \mathbf{A}_{m_1} + \mathbf{W} \mathbf{W}^\top \mathbf{A} \mathbf{W} \mathbf{W}^\top.$$

The following lemma lists some of the properties of $\tilde{\mathbf{A}}$.

Lemma 3.3. *It holds that*

1. $\dim(\text{Range}(\tilde{\mathbf{A}})) = k$ and $\dim(\text{Range}(\mathbf{W})) = k - m_1$;
2. $\mathcal{U}_{m_1} \subseteq \text{Range}(\tilde{\mathbf{A}}) \subseteq \mathcal{U}_{m_2}$ and $\text{Range}(\tilde{\mathbf{A}} - \mathbf{A}_{m_1}) \subseteq \mathcal{U}_{m_1:m_2}$, where $\mathcal{U}_{m_2} = \mathcal{U}_{m_1} \oplus \mathcal{U}_{m_1:m_2}$.
3. $\|\hat{\mathbf{U}}_k^\top \tilde{\mathbf{U}}_\perp\|_2, \|\tilde{\mathbf{U}}^\top \hat{\mathbf{U}}_{n-k}\|_2 \leq 2\epsilon$, where $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{U}}_\perp$ are orthonormal basis of $\text{Range}(\tilde{\mathbf{A}})$ and $\text{Null}(\tilde{\mathbf{A}})$, respectively.

Proof. Properties 1 and 2 are obviously true by the definition of \mathcal{W} and $\tilde{\mathbf{A}}$. For property 3, note that both $\|\hat{\mathbf{U}}_k^\top \tilde{\mathbf{U}}_\perp\|_2$ and $\|\tilde{\mathbf{U}}^\top \hat{\mathbf{U}}_{n-k}\|_2$ are equal to $\sin \angle(\tilde{\mathcal{U}}, \hat{\mathcal{U}}_k)$. Hence it suffices to show that $\|\hat{\mathbf{U}}_{n-k}^\top \tilde{\mathbf{U}}\|_2 \leq 2\epsilon$. Invoking Lemmas 3.1 and 3.2 we have that $\|\hat{\mathbf{U}}_{n-k}^\top \tilde{\mathbf{U}}\|_2 \leq \|\hat{\mathbf{U}}_{n-k}^\top \mathbf{U}_{m_1}\|_2 + \|\hat{\mathbf{U}}_{n-k}^\top \mathbf{W}\|_2 \leq \epsilon + \epsilon = 2\epsilon$. \square

Decompose $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F$ as

$$\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq \|\mathbf{A} - \tilde{\mathbf{A}}\|_F + \|\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}\|_F \leq \|\mathbf{A} - \tilde{\mathbf{A}}\|_F + \sqrt{2k} \|\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}\|_2. \quad (10)$$

Here the last inequality holds because both $\hat{\mathbf{A}}_k$ and $\tilde{\mathbf{A}}$ have rank at most k . Lemmas 3.4 and 3.5 give separate upper bounds for $\|\mathbf{A} - \tilde{\mathbf{A}}\|_F$ and $\|\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}\|_2$.

Lemma 3.4. *If $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})^2$ for $\epsilon \in (0, 1/4]$ then $\|\mathbf{A} - \tilde{\mathbf{A}}\|_F \leq (1 + 32\epsilon) \|\mathbf{A} - \mathbf{A}_k\|_F$.*

Proof. Let $\mathcal{U}_{m_1:m_2}$ be the $(m_2 - m_1)$ -dimensional linear subspace such that $\mathcal{U}_{m_2} = \mathcal{U}_{m_1} \oplus \mathcal{U}_{m_1:m_2}$. Define $\mathbf{A}_{m_1:m_2} = \mathbf{U}_{m_1:m_2} \boldsymbol{\Sigma}_{m_1:m_2} \mathbf{U}_{m_1:m_2}^\top$, where $\boldsymbol{\Sigma}_{m_1:m_2} = \text{diag}(\sigma_{m_1+1}(\mathbf{A}), \dots, \sigma_{m_2}(\mathbf{A}))$ and $\mathbf{U}_{m_1:m_2}$ is an orthonormal basis associated with $\mathcal{U}_{m_1:m_2}$. We then have

$$\begin{aligned} \|\mathbf{A} - \tilde{\mathbf{A}}\|_F^2 &= \|\mathbf{A}_{n-m_1} - \mathbf{W}\mathbf{W}^\top \mathbf{A}\mathbf{W}\mathbf{W}^\top\|_F^2 \\ &\stackrel{(a)}{=} \|\mathbf{A}_{n-m_2}\|_F^2 + \|\mathbf{A}_{m_1:m_2} - \mathbf{W}\mathbf{W}^\top \mathbf{A}\mathbf{W}\mathbf{W}^\top\|_F^2 \\ &\stackrel{(b)}{=} \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + \|\mathbf{A}_{m_1:m_2} - \mathbf{W}\mathbf{W}^\top \mathbf{A}_{m_1:m_2} \mathbf{W}\mathbf{W}^\top\|_F^2 \\ &\stackrel{(c)}{=} \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + \|\mathbf{A}_{m_1:m_2}\|_F^2 - \|\mathbf{W}\mathbf{W}^\top \mathbf{A}_{m_1:m_2} \mathbf{W}\mathbf{W}^\top\|_F^2. \end{aligned}$$

Here in (a) we apply $\text{Range}(\tilde{\mathbf{A}} - \mathbf{A}_{m_1}) \subseteq \mathcal{U}_{m_1:m_2}$ and the Pythagorean theorem (Lemma B.2) with $\mathbf{P} = \mathbf{U}_{m_1:m_2}$, in (b) we apply $\mathcal{W} \subseteq \mathcal{U}_{m_1:m_2}$, and in (c) we apply the Pythagorean theorem again with $\mathbf{P} = \mathbf{W}$. Note that $\|\mathbf{W}\mathbf{W}^\top \mathbf{A}_{m_1:m_2} \mathbf{W}\mathbf{W}^\top\|_F^2 = \|\mathbf{W}^\top \mathbf{A}_{m_1:m_2} \mathbf{W}\|_F^2$. Applying Poincaré separation theorem (Lemma B.3) where $\mathbf{X} = \boldsymbol{\Sigma}_{m_1:m_2}$ and $\mathbf{P} = \mathbf{U}_{m_1:m_2}^\top \mathbf{W}$, we have $\|\mathbf{W}^\top \mathbf{A}_{m_1:m_2} \mathbf{W}\|_F^2 \geq \sum_{j=m_2-k+1}^{m_2-m_1} \sigma_j(\mathbf{A}_{m_1:m_2})^2 = \sum_{j=m_1+m_2-k+1}^{m_2} \sigma_j(\mathbf{A})^2$. Subsequently,

$$\begin{aligned} \|\mathbf{A} - \tilde{\mathbf{A}}\|_F^2 &\leq \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + \sum_{j=m_1+1}^{m_1+m_2-k} \sigma_j(\mathbf{A})^2 \leq \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + (m_2 - k) \sigma_{m_1+1}(\mathbf{A})^2 \\ &\stackrel{(a')}{\leq} \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + (m_2 - k)(1 + 2\epsilon)^2 \sigma_{k+1}(\mathbf{A})^2 \\ &\stackrel{(b')}{\leq} \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + (m_2 - k) \left(\frac{1 + 2\epsilon}{1 - 2\epsilon} \right)^2 \sigma_{m_2}(\mathbf{A})^2 \\ &\stackrel{(c')}{\leq} \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + (m_2 - k) \sigma_{m_2}(\mathbf{A})^2 + 32(m_2 - k) \epsilon \sigma_{m_2}(\mathbf{A})^2 \\ &\stackrel{(d')}{\leq} (1 + 32\epsilon) \|\mathbf{A} - \mathbf{A}_k\|_F^2. \end{aligned}$$

Here in (a') we apply the definition of m_1 that $\sigma_{m_1+1} \leq (1 + 2\epsilon) \sigma_{k+1}(\mathbf{A})$, in (b') we apply the definition of m_2 that $\sigma_{m_2}(\mathbf{A}) \geq \sigma_k(\mathbf{A}) - 2\epsilon \sigma_{k+1}(\mathbf{A}) \geq (1 - 2\epsilon) \sigma_{k+1}(\mathbf{A})$, and (c') is due to the fact that $\left(\frac{1+2\epsilon}{1-2\epsilon} \right)^2 \leq 1 + 32\epsilon$ for all $\epsilon \in (0, 1/4]$. Finally, (d') holds because $(m_2 - k) \sigma_{m_2}(\mathbf{A})^2 \leq \sum_{j=k+1}^{m_2} \sigma_j(\mathbf{A})^2$ and $\|\mathbf{A} - \mathbf{A}_k\|_F^2 = \|\mathbf{A} - \mathbf{A}_{m_2}\|_F^2 + \sum_{j=k+1}^{m_2} \sigma_j(\mathbf{A})^2$. \square

Lemma 3.5. *If $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$ for $\epsilon \in (0, 1/4]$ then $\|\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}\|_2 \leq 102\epsilon^2 \|\mathbf{A} - \mathbf{A}_k\|_2$.*

Proof. Recall the definition that $\tilde{\mathcal{U}} = \text{Range}(\tilde{\mathbf{A}})$ and $\tilde{\mathcal{U}}_\perp = \text{Null}(\tilde{\mathbf{A}})$. Consider $\|\mathbf{v}\|_2 = 1$ such that $\mathbf{v}^\top (\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}) \mathbf{v} = \|\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}\|_2$. Because \mathbf{v} maximizes $\mathbf{v}^\top (\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}) \mathbf{v}$ over all unit-length vectors, it must lie in the range of $(\hat{\mathbf{A}}_k - \tilde{\mathbf{A}})$ because otherwise the component outside the range will not contribute. Therefore, we can choose \mathbf{v} that $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$ where $\mathbf{v}_1 \in \text{Range}(\hat{\mathbf{A}}_k) = \tilde{\mathcal{U}}_k$ and $\mathbf{v}_2 \in \text{Range}(\tilde{\mathbf{A}}) = \tilde{\mathcal{U}}$. Subsequently, we have that

$$\mathbf{v} = \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top \mathbf{v} + \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \hat{\mathbf{U}}_{n-k} \hat{\mathbf{U}}_{n-k}^\top \mathbf{v} \quad (11)$$

$$= \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \mathbf{v} + \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \mathbf{v}. \quad (12)$$

Consider the following decomposition:

$$\left| \mathbf{v}^\top (\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}) \mathbf{v} \right| \leq \left| \mathbf{v}^\top (\hat{\mathbf{A}} - \mathbf{A}) \mathbf{v} \right| + \left| \mathbf{v}^\top (\hat{\mathbf{A}}_k - \hat{\mathbf{A}}) \mathbf{v} \right| + \left| \mathbf{v}^\top (\mathbf{A} - \tilde{\mathbf{A}}) \mathbf{v} \right|.$$

The first term $|\mathbf{v}^\top (\hat{\mathbf{A}} - \mathbf{A})\mathbf{v}|$ is trivially upper bounded by $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$. For the second term, we have

$$\begin{aligned} |\mathbf{v}^\top (\hat{\mathbf{A}}_k - \tilde{\mathbf{A}})\mathbf{v}| &= |\mathbf{v}^\top \hat{\mathbf{U}}_{n-k} \hat{\Sigma}_{n-k} \hat{\mathbf{U}}_{n-k}^\top \mathbf{v}| \\ &\stackrel{(a)}{=} |\mathbf{v}^\top \hat{\mathbf{U}}_{n-k} \hat{\mathbf{U}}_{n-k}^\top \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \hat{\mathbf{U}}_{n-k} \hat{\Sigma}_{n-k} \hat{\mathbf{U}}_{n-k}^\top \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \hat{\mathbf{U}}_{n-k} \hat{\mathbf{U}}_{n-k}^\top \mathbf{v}| \\ &\leq \|\hat{\mathbf{U}}_{n-k}^\top \tilde{\mathbf{U}}\|_2^4 \|\hat{\mathbf{U}}_{n-k}\|_2 \stackrel{(b)}{\leq} 16\epsilon^4 \sigma_{k+1}(\hat{\mathbf{A}}) \stackrel{(c)}{\leq} 16\epsilon^4 (1 + \epsilon^2) \sigma_{k+1}(\mathbf{A}). \end{aligned}$$

Here in (a) we apply Eq. (11); in (b) we apply Property 3 of Lemma 3.3, and (c) is due to Weyl's inequality (Lemma B.4) that $\sigma_{k+1}(\hat{\mathbf{A}}) \leq \sigma_{k+1}(\mathbf{A}) + \|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq (1 + \epsilon^2) \sigma_{k+1}(\mathbf{A})$.

For the third term, note that $\tilde{\mathbf{A}} = \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \mathbf{A} \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top$ because $\text{Range}(\tilde{\mathbf{A}}) \subseteq \mathcal{U}_{m_2} \subseteq \text{Range}(\mathbf{A})$ by Lemma 3.3. Subsequently,

$$\mathbf{A} - \tilde{\mathbf{A}} = \underbrace{\tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \mathbf{A} \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top}_{\mathbf{B}_1} + \underbrace{\tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \mathbf{A} \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top}_{\mathbf{B}_2} + \underbrace{\tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \mathbf{A} \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top}_{\mathbf{B}_2^\top}.$$

It then suffices to upper bound $|\mathbf{v}^\top \mathbf{B}_1 \mathbf{v}|$ and $|\mathbf{v}^\top \mathbf{B}_2 \mathbf{v}|$ separately. For \mathbf{B}_1 we have

$$\begin{aligned} |\mathbf{v}^\top \mathbf{B}_1 \mathbf{v}| &\stackrel{(a')}{=} |\mathbf{v}^\top \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \mathbf{A} \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \mathbf{v}| \\ &\leq \|\tilde{\mathbf{U}}_\perp^\top \hat{\mathbf{U}}_k\|_2^4 \|\tilde{\mathbf{U}}_\perp^\top \mathbf{A} \tilde{\mathbf{U}}_\perp\|_2 \\ &\stackrel{(b')}{\leq} 16\epsilon^4 \|\tilde{\mathbf{U}}_\perp^\top \mathbf{A} \tilde{\mathbf{U}}_\perp\|_2 \stackrel{(c')}{\leq} 16\epsilon^4 \sigma_{m_1+1}(\mathbf{A}) \stackrel{(d')}{\leq} 16\epsilon^4 (1 + 2\epsilon) \sigma_{k+1}(\mathbf{A}). \end{aligned}$$

Here in (a') we apply Eq. (12); in (b') we apply Property 3 of Lemma 3.3; (c') follows the property that $\tilde{\mathbf{U}}_\perp \in \mathcal{U}_{n-m_1}$, and finally (d') follows from the definition of m_1 that $\sigma_{m_1+1}(\mathbf{A}) \leq (1 + 2\epsilon) \sigma_{k+1}(\mathbf{A})$.

For \mathbf{B}_2 , we have that

$$\begin{aligned} |\mathbf{v}^\top \mathbf{B}_2 \mathbf{v}| &= |\mathbf{v}^\top \tilde{\mathbf{U}} \tilde{\mathbf{U}}^\top \mathbf{A} \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top \tilde{\mathbf{U}}_\perp \tilde{\mathbf{U}}_\perp^\top \mathbf{v}| \\ &\leq \|\mathbf{A} \tilde{\mathbf{U}}_\perp\|_2 \|\tilde{\mathbf{U}}_\perp^\top \hat{\mathbf{U}}_k\|_2^2 \leq \epsilon^2 (1 + 8\epsilon) \sigma_{k+1}(\mathbf{A}). \end{aligned}$$

Combining all inequalities and noting that $\epsilon \in (0, 1/4]$, we obtain

$$\begin{aligned} \|\hat{\mathbf{A}}_k - \tilde{\mathbf{A}}\|_2 &\leq \epsilon^2 \sigma_{k+1}(\mathbf{A}) + 16\epsilon^4 (1 + 2\epsilon + \epsilon^2) \sigma_{k+1}(\mathbf{A}) + 32\epsilon^2 (1 + 8\epsilon) \sigma_{k+1}(\mathbf{A}) \\ &\leq 102\epsilon^2 \sigma_{k+1}(\mathbf{A}). \end{aligned}$$

□

4 Discussion

We mention two potential directions to further extend results of this paper.

4.1 Model selection for general high-rank matrices

The validity of Theorem 2.1 depends on the condition $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$, which could be hard to verify if $\sigma_{k+1}(\mathbf{A})$ is unknown and difficult to estimate. Furthermore, for general high-rank matrices, the *model selection* problem of determining an appropriate (or even optimal) cut-off rank k requires knowledge

of the distribution of the entire spectrum of an unknown data matrix, which is even more challenging to obtain.

One potential approach is to impose a parametric pattern of decay of the eigenvalues (e.g., polynomial and exponential decay), and to estimate a small set of parameters (e.g., degree of polynomial) from the noisy observations $\hat{\mathbf{A}}$. Afterwards, the optimal cut-off rank k could be determined by a theoretical analysis, similar to the examples in Corollaries 2.1 and 2.2. Another possibility is to use repeated sampling techniques such as bootstrap in a stochastic problem (e.g., matrix de-noising) to estimate the “bias” term $\|\mathbf{A} - \mathbf{A}_k\|_F$, as the variance term $\sqrt{k}\nu$ is known or easy to estimate.

4.2 Minimax rates for polynomial spectral decay

Consider the class of PSD matrices whose eigenvalues follow a polynomial (power-law) decay: $\Theta(\beta, n) = \{\mathbf{A} \in \mathbb{R}^{n \times n} : \mathbf{A} \succ 0, \sigma_j(\mathbf{A}) = j^{-\beta}\}$. We are interested in the following minimax rates for completing or de-noising matrices in $\Theta(\beta, n)$:

Question 1 (Completion of $\Theta(\beta, n)$). Fix $n \in \mathbb{N}$, $p \in (0, 1)$ and define $N = pn^2$. For $\mathbf{M} \in \Theta(\beta, n)$, let $\hat{\mathbf{A}}_{ij} = \mathbf{M}_{ij}$ with probability p and $\hat{\mathbf{A}}_{ij} = 0$ with probability $1 - p$. Also let $\Lambda(\mu_0, n) = \{\mathbf{M} \in \mathbb{R}^{n \times n} : n\|\mathbf{M}\|_{\max} \leq \mu_0\|\mathbf{M}\|_F\}$ be the class of all non-spiky matrices. Determine

$$R_1(\mu_0, \beta, n, N) := \inf_{\hat{\mathbf{A}} \mapsto \hat{\mathbf{M}}} \sup_{\mathbf{M} \in \Theta(\beta, n) \cap \Lambda(\mu_0, n)} \mathbb{E}\|\hat{\mathbf{M}} - \mathbf{M}\|_F^2.$$

Question 2 (De-noising of $\Theta(\beta, n)$). Fix $n \in \mathbb{N}$, $\nu > 0$ and let $\hat{\mathbf{A}} = \mathbf{M} + \nu/\sqrt{n}\mathbf{Z}$, where \mathbf{Z} is a symmetric matrices with i.i.d. standard Normal random variables on its upper triangle. Determine

$$R_2(\nu, \beta, n) := \inf_{\hat{\mathbf{A}} \mapsto \hat{\mathbf{M}}} \sup_{\mathbf{M} \in \Theta(\beta, n)} \mathbb{E}\|\hat{\mathbf{M}} - \mathbf{M}\|_F^2.$$

Compared to existing settings on matrix completion and de-noising, we believe $\Theta(\beta, n)$ is a more natural matrix class which allows for general high-rank matrices, but also imposes sufficient spectral decay conditions so that spectrum truncation algorithms result in significant benefits. Based on Corollary 2.1 and its matching lower bounds for a larger ℓ_p class [NW12], we make the following conjecture:

Conjecture 4.1. For $\beta > 1/2$ and ν not too small, we conjecture that

$$R_1(\mu_0, \beta, n, N) \asymp C(\mu_0) \cdot \left[\frac{n}{N}\right]^{\frac{2\beta-1}{2\beta}} \quad \text{and} \quad R_2(\nu, \beta, n) \asymp [\nu^2]^{\frac{2\beta-1}{2\beta}},$$

where $C(\mu_0) > 0$ is a constant that depends only on μ_0 .

References

- [AM07] Dimitris Achlioptas and Frank McSherry. Fast computation of low-rank matrix approximations. *Journal of the ACM*, 54(2):9, 2007.
- [AZL16] Zeyuan Allen-Zhu and Yuanzhi Li. Even faster svd decomposition yet without agonizing pain. In *Advances in Neural Information Processing Systems*, pages 974–982, 2016.
- [BDWY16] Maria Florina Balcan, Simon S Du, Yining Wang, and Adams Wei Yu. An improved gap-dependency analysis of the noisy power method. In *29th Annual Conference on Learning Theory*, pages 284–309, 2016.
- [BX⁺15] Florentina Bunea, Luo Xiao, et al. On the sample covariance matrix estimator of reduced effective rank population matrices, with applications to fpca. *Bernoulli*, 21(2):1200–1230, 2015.

- [CP10] Emmanuel J Candes and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [CR12] Emmanuel Candes and Benjamin Recht. Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119, 2012.
- [CRZ13] T Tony Cai, Zhao Ren, and Harrison H Zhou. Optimal rates of convergence for estimating toeplitz covariance matrices. *Probability Theory and Related Fields*, 156(1-2):101–143, 2013.
- [CRZ⁺16] T Tony Cai, Zhao Ren, Harrison H Zhou, et al. Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation. *Electronic Journal of Statistics*, 10(1):1–59, 2016.
- [CXCS16] Yudong Chen, Huan Xu, Constantine Caramanis, and Sujay Sanghavi. Matrix completion with column manipulation: Near-optimal sample-robustness-rank tradeoffs. *IEEE Transactions on Information Theory*, 62(1):503–526, 2016.
- [CZ⁺12] T Tony Cai, Harrison H Zhou, et al. Optimal rates of convergence for sparse covariance matrix estimation. *The Annals of Statistics*, 40(5):2389–2420, 2012.
- [CZZ⁺10] T Tony Cai, Cun-Hui Zhang, Harrison H Zhou, et al. Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*, 38(4):2118–2144, 2010.
- [DG⁺14] David Donoho, Matan Gavish, et al. Minimax risk of matrix denoising by singular value thresholding. *The Annals of Statistics*, 42(6):2413–2440, 2014.
- [DGM13] David L Donoho, Matan Gavish, and Andrea Montanari. The phase transition of matrix recovery from gaussian measurements matches the minimax mse of matrix denoising. *Proceedings of the National Academy of Sciences*, 110(21):8405–8410, 2013.
- [EBN12] Brian Eriksson, Laura Balzano, and Robert D Nowak. High-rank matrix completion. In *AISTATS*, pages 373–381, 2012.
- [GD14] Matan Gavish and David L Donoho. The optimal hard threshold for singular values is $4/\sqrt{3}$. *IEEE Transactions on Information Theory*, 60(8):5040–5053, 2014.
- [Har14] Moritz Hardt. Understanding alternating minimization for matrix completion. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 651–660. IEEE, 2014.
- [HW14] Moritz Hardt and Mary Wootters. Fast matrix completion without the condition number. In *COLT*, pages 638–678, 2014.
- [JN14] Prateek Jain and Praneeth Netrapalli. Fast exact matrix completion with finite samples. *arXiv preprint*, 2014.
- [JNS13] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674. ACM, 2013.
- [KLT11] Vladimir Koltchinskii, Karim Lounici, and Alexandre B Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, pages 2302–2329, 2011.
- [KMO10] Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.

- [KS11] Alois Kneip and Pascal Sarda. Factor models and variable selection in high-dimensional regression analysis. *The Annals of Statistics*, pages 2410–2447, 2011.
- [LWS15] Ziqi Liu, Yu-Xiang Wang, and Alexander Smola. Fast differentially private matrix factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 171–178. ACM, 2015.
- [MJC⁺14] Lester Mackey, Michael I Jordan, Richard Y Chen, Brendan Farrell, Joel A Tropp, et al. Matrix concentration inequalities via the method of exchangeable pairs. *The Annals of Probability*, 42(3):906–945, 2014.
- [NW12] Sahand Negahban and Martin J Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *The Journal of Machine Learning Research*, 13(1):1665–1697, 2012.
- [SL16] Ruoyu Sun and Zhi-Quan Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Transactions on Information Theory*, 62(11):6535–6579, 2016.
- [VdV00] Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- [ZYJZ15] Lijun Zhang, Tianbao Yang, Rong Jin, and Zhi-Hua Zhou. Analysis of nuclear norm regularization for full-rank matrix completion. *arXiv preprint arXiv:1504.06817*, 2015.

A Proof of corollaries

Proof. of Corollary 2.1. We first verify the condition that $\delta \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$ for $\epsilon = 1/4$ and the particular choice of k . Because $k \leq \lfloor C_1 \delta^{-1/\beta} \rfloor - 1$, we have that $\sigma_{k+1}(\mathbf{A}) \geq (C_1 \delta^{-1/\beta})^{-\beta}$. By carefully chosen C_1 (depending on β) the inequality $\sigma_{k+1}(\mathbf{A}) \geq \delta/16$ holds.

If $k = n - 1$ then by Theorem 2.1, $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq O(\sqrt{n} \cdot n^{-\beta}) = O(n^{-\frac{2\beta-1}{2\beta}})$. In the rest of the proof we assume $k = \lfloor C_1 \delta^{-1/\beta} \rfloor - 1$. We then have

$$\|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_{j=k+1}^n \sigma_j(\mathbf{A})^2} = \sqrt{\sum_{j=k+1}^n j^{-2\beta}} \leq \sqrt{\int_k^\infty x^{-2\beta} dx} = \sqrt{\frac{k^{-(2\beta-1)}}{2\beta-1}} \leq C(\beta) \delta^{\frac{2\beta-1}{2\beta}}.$$

Here $C(\beta) > 0$ is a constant that only depends on β . In addition,

$$\sqrt{k} \|\mathbf{A} - \mathbf{A}_k\|_2 \leq \sqrt{k} \cdot k^{-\beta} = k^{-(\beta-1/2)} \leq \tilde{C}(\beta) \delta^{\frac{2\beta}{2\beta-1}}.$$

Applying Theorem 2.1 we complete the proof of Corollary 2.1. \square

Proof. of Corollary 2.2 We first verify the condition that $\delta \leq \epsilon^2 \sigma_{k+1}(\mathbf{A})$ for $\epsilon = 1/4$ and the particular choice of k . Because $k \leq \lfloor c^{-1} \log(1/\delta) - c^{-1} \log \log(1/\delta) \rfloor - 1$, we have that $\sigma_{k+1}(\mathbf{A}) \geq \delta \log(1/\delta)$. Hence, for $\delta \in (0, e^{-16})$ it holds that $\sigma_{k+1}(\mathbf{A}) \geq \delta/16$.

If $k = n - 1$ then by Theorem 2.1, $\|\hat{\mathbf{A}}_k - \mathbf{A}\|_F \leq O(\sqrt{n} \cdot \exp\{-cn\})$. In the rest of the proof we assume $k = \lfloor C_2 \log(1/\delta) \rfloor - 1$. We then have

$$\|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_{j=k+1}^n \sigma_j(\mathbf{A})^2} = \sqrt{\sum_{j=k+1}^n \exp\{-2cj\}} \leq \sqrt{\frac{\exp\{-2ck\}}{1 - e^{-2c}}} \leq C(c) \delta \log(1/\delta),$$

where $C(c) > 0$ is a constant that only depends on c . In addition,

$$\sqrt{k} \|\mathbf{A} - \mathbf{A}_k\|_2 \leq \sqrt{k} \cdot \exp\{-ck\} \leq \delta \log(1/\delta) \cdot \sqrt{c^{-1} \log(1/\delta)} \leq \tilde{C}(c) \delta \sqrt{\log(1/\delta)^3}.$$

Applying Theorem 2.1 we complete the proof of Corollary 2.2. \square

B Technical lemmas

Lemma B.1 (Asymmetric Davis-Kahan inequality). *Fix $i \leq j \leq n$ and suppose \mathbf{X}, \mathbf{Y} are symmetric $n \times n$ matrices, with eigen-decomposition $\mathbf{X} = \mathbf{P}_i \mathbf{\Lambda}_i \mathbf{P}_i^\top + \mathbf{P}_{n-i} \mathbf{\Lambda}_{n-i} \mathbf{P}_{n-i}^\top$ and $\mathbf{Y} = \mathbf{Q}_j \mathbf{\Xi}_j \mathbf{Q}_j^\top + \mathbf{Q}_{n-j} \mathbf{\Xi}_{n-j} \mathbf{Q}_{n-j}^\top$. If $\sigma_i(\mathbf{X}) > \sigma_{j+1}(\mathbf{Y})$ then*

$$\|\mathbf{Q}_{n-j}^\top \mathbf{P}_i\|_2 \leq \frac{\|\mathbf{X} - \mathbf{Y}\|_2}{\sigma_i(\mathbf{X}) - \sigma_{j+1}(\mathbf{Y})}.$$

Proof. Consider

$$\|\mathbf{Q}_{n-j}^\top (\mathbf{X} - \mathbf{Y}) \mathbf{P}_i\|_2 = \|\mathbf{Q}_{n-j}^\top \mathbf{P}_i \mathbf{\Lambda}_i - \mathbf{\Xi}_{n-j} \mathbf{Q}_{n-j}^\top \mathbf{P}_i\|_2 \geq \|\mathbf{Q}_{n-j}^\top \mathbf{P}_i\|_2 (\sigma_i(\mathbf{X}) - \sigma_{j+1}(\mathbf{Y})).$$

Because $\sigma_i(\mathbf{X}) > \sigma_{j+1}(\mathbf{Y})$, we have that

$$\|\mathbf{Q}_{n-j}^\top \mathbf{P}_i\|_2 \leq \frac{\|\mathbf{Q}_{n-j}^\top (\mathbf{X} - \mathbf{Y}) \mathbf{P}_i\|_2}{\sigma_i(\mathbf{X}) - \sigma_{j+1}(\mathbf{Y})} \leq \frac{\|\mathbf{X} - \mathbf{Y}\|_2}{\sigma_i(\mathbf{X}) - \sigma_{j+1}(\mathbf{Y})}.$$

□

Lemma B.2 (Pythagorean theorem). *Fix $n \geq m$. Suppose \mathbf{X} is a symmetric $n \times n$ matrix and \mathbf{P} is an $n \times m$ matrix satisfying $\mathbf{P}^\top \mathbf{P} = \mathbf{I}$. Then $\|\mathbf{X}\|_F^2 = \|\mathbf{X} - \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top\|_F^2 + \|\mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top\|_F^2$.*

Proof. Expanding $\|\mathbf{X}\|_F^2$ we have that

$$\begin{aligned} \|\mathbf{X}\|_F^2 &= \|(\mathbf{X} - \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top) + \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top\|_F^2 \\ &= \|\mathbf{X} - \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top\|_F^2 + \|\mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top\|_F^2 + 2\text{tr}[(\mathbf{X} - \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top) \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top]. \end{aligned}$$

It suffices to prove that the trace term is zero:

$$\begin{aligned} \text{tr}[(\mathbf{X} - \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top) \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top] &= \text{tr}(\mathbf{X} \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top) - \text{tr}(\mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top) \\ &\stackrel{(*)}{=} \text{tr}(\mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P}) - \text{tr}(\mathbf{P}^\top \mathbf{X} \mathbf{P} \mathbf{P}^\top \mathbf{X} \mathbf{P}) \\ &= 0. \end{aligned}$$

Here $(*)$ is due to $\mathbf{P}^\top \mathbf{P} = \mathbf{I}$.

□

Lemma B.3 (Poincaré separation theorem). *Fix $n \geq m$. Suppose \mathbf{X} is a symmetric $n \times n$ matrix, \mathbf{P} is an $n \times m$ matrix that satisfies $\mathbf{P}^\top \mathbf{P} = \mathbf{I}$, and $\mathbf{Y} = \mathbf{P}^\top \mathbf{X} \mathbf{P}$. Let $\sigma_1(\mathbf{X}) \geq \dots \geq \sigma_n(\mathbf{X})$ and $\sigma_1(\mathbf{Y}) \geq \dots \geq \sigma_m(\mathbf{Y})$ be the eigenvalues of \mathbf{X} and \mathbf{Y} in descending order. Then*

$$\sigma_i(\mathbf{X}) \geq \sigma_i(\mathbf{Y}) \geq \sigma_{n-m+i}(\mathbf{X}), \quad i = 1, \dots, m.$$

Lemma B.4 (Weyl's monotonicity theorem). *Suppose \mathbf{X}, \mathbf{Y} are $n \times n$ symmetric matrices, and let $\sigma_1(\mathbf{X}) \geq \dots \geq \sigma_n(\mathbf{X})$, $\sigma_1(\mathbf{Y}) \geq \dots \geq \sigma_n(\mathbf{Y})$ and $\sigma_1(\mathbf{X} + \mathbf{Y}) \geq \dots \geq \sigma_n(\mathbf{X} + \mathbf{Y})$ denote the eigenvalues of \mathbf{X} , \mathbf{Y} and $\mathbf{X} + \mathbf{Y}$ in descending order. Then*

$$\sigma_{i+j-1}(\mathbf{X} + \mathbf{Y}) \leq \sigma_i(\mathbf{X}) + \sigma_j(\mathbf{Y}), \quad 1 \leq i, j \leq n, i+j-1 \leq n.$$

In particular, setting $i = 1$ one obtains the commonly used Weyl's inequality: $|\sigma_j(\mathbf{X} + \mathbf{Y}) - \sigma_j(\mathbf{X})| \leq \|\mathbf{Y}\|_2$.